arXiv:1406.6332v1 [cs.SY] 24 Jun 2014

Combined Global and Local Search for the Falsification of Hybrid Systems *

Jan Kuřátko[1,2] and Stefan Ratschan[1]**

[1] Institute of Computer Science, Academy of Sciences of the Czech Republic
[2] Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic

Abstract. In this paper we solve the problem of finding a trajectory that shows that a given hybrid dynamical system with deterministic evolution leaves a given set of states considered to be safe. The algorithm combines local with global search for achieving both efficiency and global convergence. In local search, it exploits derivatives for efficient computation. Unlike other methods for falsification of hybrid systems with deterministic evolution, we do not restrict our search to trajectories of a certain bounded length but search for error trajectories of arbitrary length.

1 Introduction

In this paper we provide an algorithm that solves the problem of unbounded safety falsification of hybrid systems with deterministic evolution. This means that, given a hybrid system with deterministic evolution, and a set of initial and a set of unsafe states, we search for a trajectory of arbitrary length starting in an initial state and ending in an unsafe state.

Existing methods for falsification of hybrid systems with deterministic evolution roughly fall into the following two categories:

- Local search [1,26]: Such methods use local optimization to incrementally bring a starting trajectory closer to an error trajectory, ideally based on information on the derivative of the objective function. The advantage of local search is its relative efficiency. The disadvantage is that for convergence it needs to be started close enough to an error trajectory. At the very least it needs to start from a sequence of modes that contains an error trajectory. However, the number of sequences of modes grows exponentially with the length of the sequence which makes the search for starting trajectories for local search a difficult problem.
- Black-box global search [19,2]: Such methods search for error trajectories globally, but use black-box optimization techniques [13,24] that do not explicitly exploit the structure specific to hybrid systems (partially continuous

* This work was supported by the Czech Science Foundation (GAČR) grant number P202/12/J060 with institutional support RVO:67985807.
** ORCID: 0000-0003-1710-1513

behavior, unbounded time variable). This extends their applicability (e.g., to Simulink models), but this may also result in loss of efficiency and restrict search to trajectories up to a given user-provided length. Of course it is possible to repeatedly restart such methods with higher upper bounds on the trajectory length, but every restart loses the information computed before.

The contribution of this paper is an algorithm that combines the scalability of local search with global convergence for error trajectories of *unbounded* length. Moreover, the resulting algorithm is reasonably simple and easy to analyze and implement. Note however, that efficiency is not primary goal of this paper— since the generic structure of the resulting algorithm allows the simple incorporation more sophisticated global search techniques [13,17].

Of course, one can use local search from any result of an algorithm based on black-box global search, but this has the following drawbacks:

- It is not clear how to handle the unbounded time variable.
- Black-box global search does not explicitly exploit the structure of hybrid systems.
- Black-box global search does not exploit the fact that it is combined with a local search method and hence may both duplicate some of the efforts of local search and fail to steer its search to good starting points for local search.

Our approach is based on a standard technique in global optimization for combining local with global search, so called two-phase methods [25]. But we adapt those methods to the situation that we have here: A direct application of two-phase methods would use a search space that is spanned by variables of two kinds: the initial point of trajectories, and the trajectory length (wrt. time). However, trajectory length is special, since it is unbounded, and since computing a trajectory of the given length from a given initial point also computes all trajectories from that initial point with shorter length. Moreover, hybrid systems combine continuous with discrete behavior and local search can exploit derivatives for searching the continuous part of the states space, but no such derivatives are available for discrete search which is another obstacle to the direct application of two-phase methods.

Hence, our approach modifies two-phase methods in such a way that—instead of treating trajectory length as a problem variable—they build trajectories incrementally from trajectory segments, and use derivative based continuous local search to glue together those segments based on continuous search (the literature on numerical algorithms for solving boundary value problems calls such an approach "multiple shooting" [3,26]).

The structure of the paper is as follows: In the next section we precisely define the problem and introduce some basic definitions. In Section 3 we introduce the main algorithm. In Section 4 we present an improved, more incremental version of the algorithm. In Section 5 we describe how to do local search for error trajectories. In Section 6 we provide some termination proofs for the algorithm.

In Section 7 we present computational experiments. In Section 8 we describe related work, and in Section 9 we conclude the paper.

We thank Aditya Zutshi and Sriram Sankaranarayanan for interesting discussions on the subject of this papers.

2 Problem Formulation

In this section we introduce notation and key concepts which we use, and present the problem we try to solve.

Definition 1. *A hybrid dynamical system is a quintuple* $H = (Q, \Omega, F, G, R)$*, where*

- Q *is a finite set whose elements we call* modes*;*
- $\Omega \subseteq Q \times \mathbb{R}^n$ *(the* state space *of the hybrid system);*
- F *assigns to each mode* $q \in Q$ *a system of differential equations* $F_q(t, x, \dot{x}) = 0$*, where* $(q, x) \in \Omega$ *and* $t \in \mathbb{R}^{\geq 0}$ *is time;*
- $G \subseteq \Omega$ *(the set of* guards*);*
- $R : \Omega \mapsto \Omega$ *(the* reset function*).*

For a given $q \in Q$, we will sometimes denote by X_q the set $\{x \mid (q, x) \in \Omega\}$.

Definition 2. *A trajectory of a hybrid dynamical system* H *is a sequence of the form* $((q_1, x_1), (q_2, x_2), \ldots, (q_k, x_k))$*, where* $q_i \in Q$ *and* $x_i : [0, t_i] \mapsto X_q$ *is a continuous trajectory of the system of differential equations given by* F_{q_i}*,* $i = 1, \ldots, k$*. For all* $i \in \{1, \ldots, k-1\}$*, for all* $t \in [0, t_i)$*, not* $G(q_i, x_i(t))$*, but for the trajectory endpoints* $G(q_i, x_i(t_i))$*. Moreover, the starting points of subsequent trajectories are determined by the reset function, that is,* $R((q_i, x_i(t_i))) = (q_{i+1}, x_{i+1}(0))$*.*

We call $t_i \in \mathbb{R}^{\geq 0}$ *the length of* x_i*. Moreover, we denote by* $(q_i, x_i^s) \in \Omega$ *the starting point of a trajectory* (q_i, x_i) *and* $(q_i, x_i^e) \in \Omega$ *its endpoint.*

Now we are ready to formulate the problem of *falsification of hybrid dynamical systems.*

Problem 1. Let H be a hybrid dynamical system and Init $\subset \Omega$, Unsafe $\subset \Omega$ be two sets. The set Init is called the set of initial states and the set Unsafe is called the set of unsafe states. The problem of falsification of H is to find any trajectory $((q_1, x_1), (q_2, x_2), \ldots, (q_k, x_k))$ of H such that $(q_1, x_1^s) \in$ Init and $(q_k, x_k^e) \in$ Unsafe. Such a trajectory is called an error trajectory of H.

3 Algorithm

We now present the main algorithm. Throughout this section we will assume a given hybrid system H with set of initial states Init and set of unsafe states Unsafe.

Informally, we intend to transform Problem 1 into the minimization of a cost function. A value of this cost function will measure how far or close a sequence of points in Ω is to an error trajectory. We will minimize this cost function until we find an error trajectory.

The algorithm will maintain a finite set of points $P \subseteq \Omega$ on which it will analyze the behavior of the given hybrid system H. We call sequences of elements from P *paths*. The algorithm will do this analysis by starting numerical simulations from points in P. Such a simulation will conclude that there is a trajectory from the starting point p of a simulation to the endpoint p'. This will be stored in a relation $\rightarrow$ on P that will relate all those points p, p' in P for which simulation showed that there is a trajectory from p to p' according to H. If there is a path from an initial point to an unsafe point according to the relation $\rightarrow$, we are done. However, since this is difficult to achieve, we allow paths of points in P for which subsequent points are not in $\rightarrow$. In order to measure how far such a path is from being a trajectory we will now introduce a distance measure for points in P:

Definition 3. *Given a finite set of states $P \subseteq \Omega$ and a relation $\rightarrow \subseteq P \times P$, the* distance *$d((q,x),(q',x'))$, of states (q,x) and (q',x') in P, is*

- *0, if $(q,x) \rightarrow (q',x')$, otherwise*
- *$\|x - x'\|$, if $q = q'$, and*
- *∞, otherwise.*

Here, the symbol $\|\cdot\|$ denotes the Euclidean norm.

Note that our distance function is not symmetrical because of the relation $\rightarrow$ that, in general, is not symmetrical which corresponds to the intuition that the existence of a trajectory from p to p' does not imply the existence of a trajectory from p' to p.

We measure the difficulty of getting from an initial state of H to a given state (q,x), and from a given state (q,x) to an unsafe state as follows:

Definition 4. *For a state $(q,x) \in P$*

$$d_I\left((q,x)\right) \equiv \inf_{u \in \mathrm{Init}} d\left(u,(q,x)\right), \qquad d_U\left((q,x)\right) \equiv \inf_{u \in \mathrm{Unsafe}} d\left((q,x),u\right).$$

Now we model how close a path is to yielding an error trajectory, as follows:

Definition 5. *The* cost *$c(p_1,\ldots,p_n)$ of a path $(p_1,\ldots,p_n)$ is*

$$d_I(p_1) + \sum_{i=1}^{n-1} d(p_i,p_{i+1}) + d_U(p_n).$$

Notice that we have an error trajectory of H if the cost $c(p_1,\ldots,p_n)$ is equal to zero. In practice, we are satisfied if the distances $d_I(p_1)$ and $d_U(p_n)$ are zero and $\sum_{i=1}^{n-1} d(p_i,p_{i+1}) < \varepsilon$ for some small threshold ε.

Now we can formulate our method for falsification of hybrid dynamical systems. In a similar way as two-phase methods [25] the algorithm iterates between two phases for exploring the state space of a given hybrid dynamical system.

The first phase is local optimization. For a given path of finite cost from P we compute another path with lower cost. In this phase we employ standard techniques for continuous optimization and use gradient information based on sensitivity analysis of hybrid dynamical systems [14]. If we find a local minimum which yields an error trajectory, we are finished. However, if the minimal cost is greater than a given threshold ε, we need to proceed to phase two and explore the state space further. The reader can find more details on the first phase in Section 5.

The second phase is called global exploration. If local optimization in the first phase does not produce an error trajectory, we add additional states to the set P. There are many options for adding new states such as a random sampling, states resulting from forward and backward simulation from existing states, states suggested by more sophisticated global search techniques [13,17], and even states given by a designer of the system.

The complete Algorithm 1 follows:

Input: a set of states $P \subseteq \Omega$ and a relation $\rightarrow \subseteq P \times P$ s.t.
- $p \rightarrow p'$ implies that there is a trajectory from p to p' in H
- there is a path of points in P that has finite cost with respect to $\rightarrow$

Output: an error trajectory

while *local optimization of the path with minimal cost does not yield an error trajectory* **do**
  add a new state $r \in \Omega$ to the set P
  for *some* $p \in P$ **do**
    simulate forward from p for some time to a new state p'
    $\rightarrow := \rightarrow \cup \{(p, p')\}$
  end
  for *some* $p \in P$ **do**
    simulate backward from p for some time to a new state p'
    $\rightarrow := \rightarrow \cup \{(p', p)\}$
  end
end

Algorithm 1: Combined Global and Local Search for the Falsification of Hybrid Systems

The second requirement on the input (existence of a path of finite cost) allows us to use derivative based continuous local optimization from the beginning. For fulfilling this requirement, we observe that paths can only have infinite cost due to sub-sequent states in different modes that are not connected by the relation $\rightarrow$. We can make this more precise by the following property:

Property 1. Assume a set $P \subseteq \Omega$ and $\rightarrow \subseteq P \times P$. Let $\rightarrow_Q \subseteq P \times P$ be such that $(q, x) \rightarrow_Q (q', x')$ iff $q = q'$ or $(q, x) \rightarrow (q', x')$, and let $\rightarrow_Q^*$ be the transitive

closure of $\rightarrow_Q$. If P contains at least one initial state p and one unsafe state p' such that $p \rightarrow_Q^* p'$, then there is a path of points in P that has finite cost with respect to $\rightarrow$

The necessary elements of $\rightarrow$ can be easily formed by pairs (p, p') such that $G(p)$ and $p' = R(p)$. For example, for each pair of modes (q, q') for which there are x and x' s.t. $G(q, x)$ and $(q', x') = R(q, x)$ we could add such (q, x) and (q', x'). If H has an error trajectory then this fulfills the assumptions of the above property resulting in a path of finite cost.

The algorithm stops when we find a path whose cost is lower than some threshold ε. A concrete implementation might add another stopping criterion, for example stating a maximum number of states in P in order to ensure termination for inputs that do not feature any error trajectory.

4 Algorithmic Details

4.1 Computation of the Path of Minimal Cost

We make the following observation: In Algorithm 1, one can view the problem of computation of a path of minimal cost as a problem on weighted directed graphs: The vertices of the graph are formed by the elements of the set P and there is an edge from $p \in P$ to $p' \in P$ iff the distance $d(p, p')$ is finite. The weight of this edge is given by this value $d(p, p')$. Now the path of minimal cost is the shortest path in this graph from an initial to an unsafe state. This is a classical problem in algorithm theory with solutions such as the Floyd-Warshall algorithm.

Examining the situation more closely, we observe that our problem is neither of the all-pair shortest path, nor of the single-source shortest path kind. Instead, the paths have to start in a certain given set (call it S for source) and end in another given set (call it G for goal). This can be reduced to a problem with single vertices instead of sets by introducing two new, auxiliary vertices s and g such that s has an edge of zero cost to each element of S and such that there is an edge of zero cost from each element of G to g.

Now we are left with a single-source single-goal shortest path problem (also called point-to-point shortest path). Of course, such problems can be solved by algorithms solving the single-source shortest path problem, for example, by Dijkstra's algorithm [9]. But there are also specialized algorithms, for example, algorithm based on a bi-directional [22,4] application of Dijkstra's algorithm.

4.2 Heuristics

The algorithm can be instantiated with many heuristics resulting in special versions of Algorithm 1, for example:

- Forward version: only add initial points and only do forward simulation
- Backward version: only add unsafe points and only do backward simulation

– Complete random search version: never prolong existing simulations, only simulation from newly added points.

The algorithm also leaves open the length of the employed simulations. A simply possibility is to fix a certain length at the beginning and stick to it throughout computation.

Note that simulation might run into problems, for example due to Zeno behavior, or due to the fact that it leaves the state space of the given hybrid system. In this case we simply ignore the result of simulation and continue with the algorithm. See also more on this at the end of Section 5.

It is also possible to use information from verification tools here. Especially, one can restrict the choice of points to an abstraction computed by a verification tool [10,23].

4.3 Paths of Minimal Cost

We will now investigate the form of paths of minimal cost. For a given hybrid system H we assume the following.

1. The sets $\{x \mid (q,x) \in \text{Init}\}$ and $\{x \mid (q,x) \in \text{Unsafe}\}$ are closed and convex.
2. For all $q \in Q$ the set $\{x \mid (q,x) \in \Omega\}$ is convex.
3. For $p, p' \in P$, $p \to p'$ implies there is a trajectory from p to p' in H.
4. There is at least one path of finite cost in P with respect to $\to$.

Lemma 1. *Let H be a given hybrid system and P be a set of states such that assumptions 1.–4 hold. Let $(p_1, \ldots, p_n)$ be the path of minimal cost. Let r be a state such that $r \neq p_i$, $i = 1, \ldots, n$, and neither $r \to p_i$ nor $p_i \to r$ for any $i \in \{1, \ldots, n\}$. Then the cost of a path which is formed by either including state r in $(p_1, \ldots, p_n)$ or substituting r for any p_i's, $i \in \{1, \ldots, n\}$, in $(p_1, \ldots, p_n)$ is greater or equal to $c(p_1, \ldots, p_n)$.*

The reader can find the proof of Lemma 1 in the appendix. The consequence of this lemma can be stated like this: Assuming 1.–4. the value of $c(p_1, \ldots, p_n)$, where $(p_1, \ldots, p_n)$ is the path of minimal cost, does not depend on states $p \in P$ which are in no relation to other states in the path wrt. $\to$. In other words, for a state p_i, $i = 1, \ldots, n$, which is in no relation with other states in a path, we have $c(p_1, \ldots, p_i, \ldots, p_n) = c(p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n)$.

This is important for finding paths of lower cost. Whenever we add a new state r in Algorithm 1 we should simulate either forward in time or backward in time to create a pair of states in a relation $\to$. Solitary states do not affect the resulting value of the cost function.

Lemma 2. *Let H be a given hybrid system and P be a set of states such that assumptions 1.–4 hold. Let $(p_1, \ldots, p_n)$, $p_i \in P$, be the path of minimal cost. Let p_j, $j = 2, \ldots, n-1$, be a state such that $p_{j-1} \to p_j$ and $p_j \to p_{j+1}$. Then $d(p_{j-1}, p_{j+1}) = 0$.* □

Proof. Due to transitivity of the relation $\rightarrow$, we have $p_{j-1} \rightarrow p_{j+1}$ which gives us $d(p_{j-1}, p_{j+1}) = 0$.

Lemma 2 presents us with a choice for the application of local optimization. If a path contains such a triplet (p_{j-1}, p_j, p_{j+1}), we may either work with them as two separate hybrid trajectories or we may consider a hybrid trajectory which is formed by their connection, thus removing the intermediate state from a path. On the other hand, we also have the option to split a hybrid trajectory into shorter hybrid trajectories before passing it to the local optimizer.

In the first approach we work with shorter trajectories however the resulting optimization problem has higher dimension. The latter case may, on the other hand, cause problems because it is less numerical stable [3]. The choice depends on the system of differential equations that governs the evolution of hybrid system H.

5 Local Optimization

In Algorithm 1, after we form a path $(p_1, \ldots, p_n)$ of finite cost we try to find another path of smaller cost using local search. Therefore, we solve the minimization problem in which we seek new states $\hat{p}_1, \ldots, \hat{p}_n$ which yield a path of lower cost than $(p_1, \ldots, p_n)$. Eventually, such a path of minimal cost may correspond to an error trajectory of a hybrid system.

The efficiency of such local search can be improved by exploiting the gradient of the cost $c(p_1, \ldots, p_n)$. In this section we will develop explicit formulae for the gradient of the cost function which will allow us to use efficient off-the-shelf tools for gradient-based numerical optimization to minimize the cost function.

Without loss of generality, we will assume that for all $i \in \{1, \ldots, n\}$, $p_i \rightarrow p_{i+1}$ iff i is odd. This can be easily achieved, since, due to Lemma 1 if there are solitary states that we can remove them from $(p_1, \ldots, p_n)$ without changing the value of the cost function.

Note that in contrast to early work [26], in cases where $p_i \rightarrow p_{i+1}$, p_i and p_{i+1} are *not* restricted to be in the same mode. Moreover, points are not restricted to reside in guards of the hybrid system.

We now give explicit formulae for computation of the gradient of the cost function $c(p_1, \ldots, p_n)$. Let us start with the definition of the length of a trajectory $((q_1, x_1), \ldots, (q_k, x_k))$ and its sensitivity to the change of its initial state (q_1, x_1^s) which is essential for evaluation of the gradient of the cost $c(p_1, \ldots, p_n)$.

Definition 6. *The length of a trajectory $((q_1, x_1), \ldots, (q_k, x_k))$ is defined to be the sum $t_f = \sum_{i=1}^{k} t_i$, where t_i is the length of x_i for $i = 1, \ldots, k$.*

Definition 7. *We define a function $M : \mathbb{R} \times \Omega \mapsto \Omega$ such that for a state $(q, x) \in \Omega$ and $t \in [0, t_f]$ we have*

$$M(t, (q, x)) = (q', x') \ ,$$

where (q', x') is the end-state of the trajectory of length t whose initial state is (q, x). In cases where a reset happens at time t (which results in the trajectory of length t being non-unique), we choose the unique point before the reset.

Definition 8. *The sensitivity of a trajectory $((q_1, x_1), \ldots, (q_k, x_k))$ of H to the initial state (q_1, x_1^s) is a function $S : \mathbb{R}^{\geq 0} \mapsto \mathbb{R}^{n \times n}$ such that*

$$S(t) \equiv \frac{\partial M(t, (q, x_1^s))}{\partial x_1^s}, \quad t \in [0, t_f].$$

With this sensitivity function we can measure how the states on a hybrid trajectory are affected when we change its initial state. An important observation is that $S(0)$ is the *identity* matrix. However, for hybrid systems, the function M need not be differentiable everywhere, and so the sensitivity is not defined everywhere. Computation of the sensitivity function is subtle [14].

For our path $(p_1, \ldots, p_n)$, for certain t_i, $i = 1, 3, \ldots, n-1$, we have

$$M(t_1, p_1) = p_2, \quad M(t_3, p_3) = p_4, \quad \ldots, \quad M(t_{n-1}, p_{n-1}) = p_n \ .$$

In the sequel let us use the following notation: For any state $(q, x) \in \Omega$, we denote by $\overline{(q, x)}$ its continuous part x.

Local search adjusts the position of the initial state of each trajectory together with its length such that the cost is minimized. It uses the gradient of the cost with respect to $\overline{p_i}$ and lengths t_i for $i < n$ odd. Therefore the gradient is given by the partial derivatives $\frac{\partial c}{\partial \overline{p_i}}(p_1, \ldots, p_n)$ and $\frac{\partial c}{\partial t_i}(p_1, \ldots, p_n)$, $i < n$ odd.

We will illustrate the whole process of computing the gradient of the cost function for one particular definition of distances $d_I, d(\cdot, \cdot)$ and d_U that avoids solving another minimization problem stemming from Definition 4. Hence, we put $d_I(p_1)$ and $d_U(p_n)$ to be weighted norms to some fixed states in Init, and Unsafe respectively. This amounts to the sets Init and Unsafe being ellipsoids. We denote by $u \in \Omega$ and $v \in \Omega$ the centres of these ellipsoids and by E_I, E_U symmetric positive definite matrices which characterize the size and shape of sets Init and Unsafe. Then we consider the cost of the following special form

$$\begin{aligned} c(p_1, \ldots, p_n) &= d_I(p_1) + \sum_{i=1}^{n-2} d(p_i, p_{i+1}) + d_U(p_n) \ , \\ &= \|\overline{p_1} - \overline{u}\|_{E_I}^2 + \sum_{i \text{ even}}^{n-2} \|\overline{p_i} - \overline{p_{i+1}}\|^2 + \|\overline{p_n} - \overline{v}\|_{E_U}^2 \ . \end{aligned}$$

When we use the function $M : \mathbb{R} \times \Omega \mapsto \Omega$ from Definition 7, then the cost becomes dependent on p_i and t_i, $i = 1, 3, 5, \ldots, n-1$,

$$\begin{aligned} c(p_1, p_3, \ldots, p_{n-1}, t_1, t_3, \ldots, t_{n-1}) = \|\overline{p_1} - \overline{u}\|_{E_I}^2 + \ldots \\ + \sum_{i \text{ even}}^{n-2} \|\overline{M(t_{i-1}, p_{i-1})} - \overline{p_{i+1}}\|^2 + \ldots \\ + \|\overline{M(t_{n-1}, p_{n-1})} - \overline{v}\|_{E_U}^2 \ . \end{aligned}$$

We can compute the gradient of the cost $c(p_1, p_3, \ldots, p_{n-1}, t_1, t_3, \ldots, t_{n-1})$ which consists of the following partial derivatives

$$\frac{\partial c}{\partial \overline{p_1}} = 2\left[(\overline{p_1} - \overline{u})^T E_I + \left(\overline{M(t_1, p_1)} - \overline{p_3}\right)^T \frac{\partial M}{\partial \overline{p_1}}(t_1, p_1)\right]$$

and for odd i with $1 < i < n-1$

$$\frac{\partial c}{\partial \overline{p_i}} = -2\left[\left(\overline{M(t_{i-2}, p_{i-2})} - \overline{p_i}\right)^T - \left(\overline{M(t_i, p_i)} - \overline{p_{i+2}}\right)^T \frac{\partial M}{\partial \overline{p_i}}(t_i, p_i)\right]$$

with

$$\begin{aligned}\frac{\partial c}{\partial \overline{p_{n-1}}} = -2\Bigg[&\left(\overline{M(t_{n-2}, p_{n-2})} - \overline{p_{n-1}}\right)^T - \\ &- \left(\overline{M(t_{n-1}, p_{n-1})} - \overline{v}\right)^T E_U \frac{\partial M}{\partial \overline{p}_{n-1}}(t_{n-1}, p_{n-1})\Bigg] \ .\end{aligned}$$

For odd $i < n-1$ we put

$$\frac{\partial c}{\partial t_i} = 2(\overline{M(t_i, p_i)} - \overline{p_{i+1}})^T \frac{\partial M}{\partial t_i}(t_i, p_i)$$

and

$$\frac{\partial c}{\partial t_{n-1}} = 2(\overline{M(t_{n-1}, p_{n-1})} - \overline{v})^T E_U \frac{\partial M}{\partial t_{n-1}}(t_{n-1}, p_{n-1}) \ .$$

In addition we may introduce weights into the cost function to scale the problem.

Now we can use numerical optimization algorithms with the cost function c and the gradients defined as above, to do local search for paths of minimal cost. If started close enough to an error trajectory, and if the hybrid system is sufficiently well-behaved around the error trajectory, such local search will converge (usually quickly). However, if this is not the case, local search may fail, due to various reasons:

- It may run in a local minimum that is not an error trajectory.
- There may be problems due to the fact that the sensitivity is not everywhere continuously differentiable. This corresponds to the situation where a trajectory is tangential to the boundary of a guard [14].
- Well-known problems with simulation [18] of hybrid systems might arise. For example, the simulation might run into Zeno behavior, or the ODE solver is unable to start close to the boundary of a guard.
- Optimization may result in trajectories that leave the state space of the hybrid system.

In all such cases, we simply terminate local optimization and continue with the global phase of the main algorithm.

6 Termination Proof

We assume a hybrid system H with the following properties:

- The state space of H is compact.
- There exists an error trajectory E with final point in the interior of the set of unsafe points, and an $\varepsilon > 0$ such that starting local search from any sequence of hybrid trajectories with cost not bigger than ε converges to an error trajectory.
- There exists a tube around E such that trajectories starting in this tube depend continuously on their initial value (note that for ODEs this can be ensured by Lipschitz continuous right-hand sides).

Moreover, we will study a variant of the algorithm with the following properties:

- The algorithm does at least one forward simulation in each cycle, always of length (in time) T.
- The algorithm chooses the starting point for its simulations randomly using a distribution that is non-zero on the whole state space.
- If a simulation hits an unsafe state, it finishes (so, in such cases, the length of the simulation may be shorter than T).
- The simulations are exact, that is, we ignore rounding and discretization errors of ODE solvers.

Note that the assumptions are asymmetrical wrt. time and set of initial vs. unsafe states. This is necessary since simulations have to be done in a certain direction, and since convergence requires simulations to be stopped if reaching an unsafe state.

While it is obvious that such an algorithm will densely fill the state space of the hybrid system with initial values of simulation, it is not obvious that this will eventually result in a path of small enough cost, since—from a given initial value—the trajectories follow the dynamics of the hybrid system H. Still, we have:

Theorem 1. *Under the assumption above, the algorithm finds an error trajectory with probability* 1*.*

The reader can find the proof in the appendix. Clearly one can easily get a dual version of the theorem and proof by turning around the time axis, switching initial and unsafe states etc.

Only slightly changing the proof, one can prove the following non-probabilistic version of the theorem:

Theorem 2. *Take the same assumptions as the previous theorem with the exception of choice of starting points of simulations. Instead of a choice according to some probability distribution, assume a choice of those starting points that fulfills the following property: For each $\varepsilon > 0$, there is an integer k such that for every ε-ball with center in the state space contains a simulation starting point that the algorithm has chosen in the first k iterations. Then algorithm always finds an error trajectory.*

7 Computational Experiments

Recall that one of the main goals of our method was to handle the absence of an a-priori upper bound on the length of error trajectories. In order to study the cost of having to work without this information we compare our approach (that we will call "unbounded method") with another approach that also combines global with local search, but that does simulations of fixed length (we will call it "bounded method"). The bounded method will also use derivative-based local optimization, but for initializing local optimization it randomly generates initial states in the mode containing I and simulates for the time interval $[0, T]$. Whenever the resulting trajectory reaches the mode containing the set of unsafe states U, we take it as a starting trajectory for local search for an error trajectory. If we obtain an error trajectory then we stop. Otherwise we proceed until we generate a certain number of trajectories (we will denote this number by M).

Note that any method that inspects the given hybrid system only up to a fixed time bound T, if T is too small, it will not find any error trajectory at all. The bounded method that we use here, for T too small, may never reach a mode containing U, preventing it from finding any error trajectory. Moreover, examples for which trajectories leading to the mode containing U lead over a very small guard, become arbitrarily difficult for the bounded method. So we can already conclude now—without running any experiments—that the unbounded method is superior in certain cases.

Still we do some experiments with a widely known benchmark, the Navigation benchmark with 16 modes [12]. We consider the linear dynamics $\dot{x} = Ax - Bu(i, j)$, with A and B as usual for the navigation benchmark, and

$$u(i,j) = \begin{bmatrix} \sin\left(\frac{\pi\, C(i,j)}{4}\right) \\ \cos\left(\frac{\pi\, C(i,j)}{4}\right) \end{bmatrix}, \qquad C = \begin{bmatrix} 4\ 3\ 3\ 4 \\ 4\ 4\ 4\ 4 \\ 4\ 6\ 6\ 4 \\ 1\ 0\ 7\ 6 \end{bmatrix}.$$

Assume the sets of initial and unsafe states to be *ellipsoids* such that their principal axes have length 0.2, 0.2, 2 and 2, however, I is centred at $\begin{bmatrix}0.5\ 3.5\ 0\ 0\end{bmatrix}^T$ and U is centred at $\begin{bmatrix}3.5\ 1.5\ 0\ 0\end{bmatrix}^T$. Our objective is to find any trajectory which starts in set I and reaches set U.

For our experiments we use an instantiation of the unbounded method that fulfills the requirement of starting with a set P that has a path of finite cost as follows: We initialize the set P by putting a point on each boundary of two neighboring modes, simulating forward and backward from each such point (0.05 time units in each direction), and adding the endpoints to the set P.

In the main algorithm, we add a random state to each mode (with velocities x_3 and x_4 ranging from -1 to 1) and then we simulate forward and backward in time from such a state (0.5 time units in each direction). The extremities of the resulting error trajectory (its initial and end states) are stored in P and used for obtaining a path of the minimal cost for local search. If local search

returns an error trajectory, then we stop. As in the bounded method we restrict computation, but this time, to add up to M states to the set P.

In all our experiments we do local search using the Scilab function for gradient-based numerical optimization, computing the gradient as described in Section 5. We weighted each distance between two consecutive segments by the weight $\omega = 500$ in order to prefer the continuity of resulting trajectories. We use the Scilab function *rand* for generating random states. For reducing dependence of the result on the random number generator, we always carry out 100 experiments, initializing the random number generator with a different seed (concretely, *rand("seed", i)*, where $i = 1, \ldots 100$). In all our experiments we use the value 500 for the constant M. The results are listed in Tab. 1. The column "successful falsification" lists the number of experiments (from 100) for which the method found an error trajectory. The column "average total simulation time" is the average of the length of all simulation done during a given experiment, but *only* for those experiments that succeeded in finding an error trajectory.

	successful falsification	average total simulation time
unbounded method	99	1937
bounded method, $T = 10$	85	1260
bounded method, $T = 20$	89	2935

Table 1. Computation Results

The choices $T = 10$ and $T = 20$ that we used are big enough, so the bounded method does find error trajectories, but still the success rate is lower than with the unbounded method, that does not need any bound T at all. In those cases where the bounded method actually finds an error trajectory, if the choice of T is small but large enough to reach a mode containing U, it needs less simulation than the unbounded method. But very quickly, when not choosing T small enough, also the cost of simulations increases beyond the unbounded method.

To sum up, the unbounded method significantly increases the chance of finding an error trajectory, and moreover, it also decreases the amount of simulation needed for that, except for cases where a very good bound on the error trajectory length is available.

8 Related Work

Our algorithm can be viewed as an adaptation of the Best Start two-phase method for global optimization [25] to our context.

The falsification problem can also be viewed as a boundary value problem which is a classical topic in numerical mathematics [3]. However, classical numerical methods assume a fixed final time, whereas we search for error trajectories of arbitrary length. Moreover, classical methods for boundary values problems are restricted to purely continuous systems and the formulation of boundary conditions as equalities.

Zuthsi and co-authors [26] present a method for falsification of hybrid systems that also uses multiple shooting based local search. However, the method assumes a given upper bound on the length of the error trajectory the method searches for. Moreover, their local search method always follows a given sequence of modes and transitions. They propose to search for such a sequence using tools that compute abstractions of hybrid systems, or by random search. The form of the used trajectory segments is more restricted than in our method since trajectory segments always stay in one mode, and end in the guard leading to another mode.

Abbas [1] show how to use local search for falsification of hybrid systems with affine dynamics. They propose to start the method from the result of global search algorithms [19].

The usage of abstractions for guiding local search for error trajectories has been proposed earlier [23], in combination with the usage of derivative-free algorithms for local search.

There is more related work for systems that—different from our case—allow input or have non-deterministic dynamics. In the completely discrete case this amounts to finding shortest paths in graphs [4]. We use shortest path algorithms as a sub-algorithm to find starting points for local search. Similar problems are studied in more structured domains by the field of planning [16], and in formal verification by directed model checking [11].

In the continuous case, the classical field studying algorithm for finding paths of dynamical systems that are in some sense optimal (e.g., as short as possible), is optimal control [5,6]. In recent years, also the field of planning has started to study continuous dynamical systems [16, Chapter IV: Planning Under Differential Constraints] from a different perspective. More recently, such techniques have also been applied to hybrid systems [7,8,20,21]. Planning-based techniques search globally, and do not require an upper bound on trajectory length, but they do not incorporate derivative-based local search. The only exception that we are aware of [15] uses optimal control to the result of planning in a purely sequential way, without any iteration between the two phases.

9 Conclusion

We presented an algorithm for the falsification of hybrid system that combines scalability due to local search with convergence due to global search. In future work, we will improve the algorithm in analogy to advanced two-phase methods [25], such as clustering methods that exploit the regions of attraction to local optima of the used local search technique.

References

1. H. Abbas and G. Fainekos. Linear hybrid system falsification with descent. Technical Report arXiv:1105.1733, 2011.

2. Y. Annpureddy, C. Liu, G. Fainekos, and S. Sankaranarayanan. S-TaLiRo: A tool for temporal logic falsification for hybrid systems. In P. A. Abdulla and K. M. Leino, editors, *TACAS*, volume 6605 of *LNCS*, pages 254–257. Springer Berlin Heidelberg, 2011.
3. U. M. Ascher, R. M. M. Mattheij, and R. D. Russell. *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*. SIAM, 1995.
4. D. P. Bertsekas. *Network optimization: continuous and discrete models*. Athena Scientific Belmont, 1998.
5. J. T. Betts. Survey of numerical methods for trajectory optimization. *Journal of Guidance, Control, and Dynamics*, 21(2), 1998.
6. M. S. Branicky, V. S. Borkar, and S. K. Mitter. A unified framework for hybrid control: Model and optimal control theory. *IEEE Transactions on Automatic Control*, 43(1):31–45, 1998.
7. M. S. Branicky, M. M. Curtiss, J. Levine, and S. Morgan. Sampling-based planning, control and verification of hybrid systems. *IEE Proceedings-Control Theory and Applications*, 153(5):575–590, 2006.
8. T. Dang and T. Nahhal. Coverage-guided test generation for continuous and hybrid systems. *Formal Methods in System Design*, 34(2):183–213, 2009.
9. E. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.
10. T. Dzetkulič and S. Ratschan. Incremental computation of succinct abstractions for hybrid systems. In *FORMATS 2011*, volume 6919 of *LNCS*, pages 271–285. Springer, Heidelberg (2011), 2011.
11. S. Edelkamp, V. Schuppan, D. Bonaki, A. Wijs, A. Fehnker, and H. Aljazzar. Survey on directed model checking. In D. Peled and M. Wooldridge, editors, *Model Checking and Artificial Intelligence*, volume 5348 of *Lecture Notes in Computer Science*, pages 65–89. Springer Berlin Heidelberg, 2009.
12. A. Fehnker and F. Ivančić. Benchmarks for hybrid systems verification. In R. Alur and G. J. Pappas, editors, *HSCC'04*, number 2993 in LNCS. Springer, 2004.
13. M. Gendreau and J.-Y. Potvin, editors. *Handbook of Metaheuristics*. Springer, 2nd edition, 2010.
14. I. Hiskens and M. Pai. Trajectory sensitivity analysis of hybrid systems. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 47(2):204–220, Feb 2000.
15. F. Lamiraux, E. Ferré, and E. Vallée. Kinodynamic motion planning: connecting exploration trees using trajectory optimization methods. In *IEEE International Conference on Robotics and Automation*, volume 4, pages 3987–3992. IEEE, 2004.
16. S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
17. M. Locatelli and F. Schoen. *Global Optimization–Theory, Algorithms, and Applications*. SIAM, 2013.
18. P. J. Mosterman. An overview of hybrid simulation phenomena and their support by simulation packages. In F. Vaandrager and J. van Schuppen, editors, *HSCC'99*, number 1569 in LNCS. Springer, 1999.
19. T. Nghiem, S. Sankaranarayanan, G. Fainekos, F. Ivančić, A. Gupta, and G. J. Pappas. Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In *HSCC '10*, pages 211–220. ACM, 2010.
20. E. Plaku, L. E. Kavraki, and M. Y. Vardi. Falsification of LTL safety properties in hybrid systems. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 5505 of *LNCS*, pages 368–382, 2009.

21. E. Plaku, L. E. Kavraki, and M. Y. Vardi. Hybrid systems: from verification to falsification by combining motion planning and discrete search. *Formal Methods in System Design*, 34(2):157–182, APR 2009.
22. I. Pohl. Bi-directional search. *Machine Intelligence*, 6:124–140, 1971.
23. S. Ratschan and J.-G. Smaus. Finding errors of hybrid systems by optimising an abstraction-based quality estimate. In C. Dubois, editor, *Tests and Proofs*, volume 5668 of *LNCS*, pages 153–168. Springer, 2009.
24. L. M. Rios and N. V. Sahinidis. Derivative-free optimization: A review of algorithms and comparison of software implementations. *Journal of Global Optimization*, pages 1–47, 2012.
25. F. Schoen. Two-phase methods for global optimization. In P. Pardalos and H. Romeijn, editors, *Handbook of Global Optimization*, volume 62 of *Nonconvex Optimization and Its Applications*, pages 151–177. Springer US, 2002.
26. A. Zutshi, S. Sankaranarayanan, J. V. Deshmukh, and J. Kapinski. A trajectory splicing approach to concretizing counterexamples for hybrid systems. In *CDC'13*, 2013.

A Appendix

Proof of Lemma 1:

Proof. The trivial case is when r does not share a mode with any p_i, $i = 1, \ldots, n$. Then it follows from Definition 3 that $d(r, p_i) = d(p_i, r) = \infty$ for all $i = 1, \ldots, n$. Since we assume $c(p_1, \ldots, p_n)$ is finite we get $d_I(r) = d_U(r) = \infty$ as well. Therefore, such a new state r cannot improve the value of $c(p_1, \ldots, p_n)$.

In order to improve on the value of $c(p_1, \ldots, p_n)$ state r must share a mode with some p_i's in $(p_1, \ldots, p_n)$. These p_i's form a subpath $(p_1^s, \ldots, p_k^s)$ and suppose states p_i^s do not belong to modes where Init and Unsafe are.

First, assume we have a new path $(p_1^s, \ldots, p_i^s, r, p_{i+1}^s, \ldots, p_k^s)$. Then, it is clear that because of the triangle inequality $c(p_1^s, \ldots, p_i^s, r, p_{i+1}^s, \ldots, p_k^s) \geq c(p_1^s, \ldots, p_k^s)$. Hence we may want to substitute state r for one or more states p_i^s. This is a more complicated, however, the same argument with the triangle inequality together with the minimality of $c(p_1, \ldots, p_n)$ apply. Note that, when we use r instead of some p_i^s, we may remove a state which is in relation given by $\rightarrow$. If this is a case, the resulting cost gets larger, even infinite, and not lower as we intended.

Now, assume state r is in a mode where either Init or Unsafe are. Since, this turns to be a symmetrical problem, we shall investigate only the case when r and Init share the same mode. If $r \in$ Init, then we are done, since $d_I(p_1) = \min_{u \in \text{Init}} d(u, p_1)$. On the other hand, if $r \notin$ Init, we use the following argument. Since Init is closed, bounded and convex, there exists an unique element $p_1^\star \in$ Init such that

$$d(p_1^\star, p_1) = \min_{u \in \text{Init}} d(u, p_1).$$

Again, using the triangle inequality $d(p_1^\star, p_1) \leq d(p_1^\star, r) + d(r, p_1)$. In addition, according to the argument above, we cannot use r instead of any state from path $(p_1, \ldots, p_n)$. □

Proof of Theorem 1:

Proof. Let E be the error trajectory, and let L be its length wrt. time. Without loss of generality we assume that $L \geq T$.

We will work with a sequence of hybrid trajectories with $\lceil \frac{L}{T} \rceil$ segments. Observe that—according to Definition 5—the allowed cost ε from which convergence of local search is guaranteed, consists of $\lceil \frac{L}{T} \rceil + 1$ summands ($\lceil \frac{L}{T} \rceil - 1$ times the cost for joining segments, the cost to initial, and the cost to unsafe). We will distribute the allowed cost ε equally among those, allowing error $\varepsilon' := \frac{\varepsilon}{\lceil \frac{L}{T} \rceil} + 1$ for each summand.

We will now analyze the summands of the cost:

- The cost resulting from the starting point of the simulation: With probability 1 the algorithm will eventually produce a segment
  - the starting point of the segment is closer than ε' from the starting point of E, and for which
  - all other points have distance less than $\frac{\varepsilon'}{2}$ from a point on E (due to continuous dependency of trajectories on the initial values).
- The cost between the end-point of any segment and starting point of the next segment: We first consider the first segment: Let p be the point on E that is closest to the end-point of the first segment. This distance is smaller than $\frac{\varepsilon'}{2}$ since the algorithm chose the starting point of the segment close enough to the starting point of E. With probability 1 the algorithm will eventually start another simulation such that
  - the starting point of the segment is closer than $\frac{\varepsilon'}{2}$ from p, and for which
  - all other points have distance less than $\frac{\varepsilon'}{2}$ from a point on E (due to continuous dependency of trajectories on the initial values).

  We can continue with the same argument along E until the final segment.
- The cost between the end-point of the final segment and the set of unsafe states: Let p be the point on E that is closest to the end-point of the last but one segment, as in the previous item. Since the final point of E is in the interior of the set of unsafe points, there is a $\mu > 0$ such that all points of distance smaller than μ from this final point is also unsafe. Again, with probability 1 we will eventually start another simulation whose distance from E is smaller than μ which means that this simulation will hit an unsafe state.

The resulting path has cost not bigger than ε. Hence local search from this path will converge to an error trajectory. The path of minimal cost used by the algorithm might be a different one that this one. However, that path of minimal cost will again have cost not bigger than ε which guarantees convergence. □